UNITED STATES DISTRICT COURT
DISTRICT OF MASSACHUSETTS

SINGULAR COMPUTING LLC, Plaintiff, v. GOOGLE LLC, Defendant.	Civil Action No. 1:19-cv-12551-FDS Hon. F. Dennis Saylor IV

DECLARATION OF SUNIL P. KHATRI, Ph.D.

TABLE OF CONTENTS

PAGE

I. EXPERTISE 1
A. Research and Professional Experience 1
B. Education 3
C. Publications, Awards, Grants, and Summary 4

II. OPINION 6

Declaration of Sunil P. Khatri, Ph.D.

I, SUNIL P. KHATRI, HEREBY DECLARE:

1. I have been retained by counsel for Patent Owner, Singular Computing LLC ("Singular"), in connection with the above-captioned litigation in the District of Massachusetts.

2. I submit this declaration in the above-captioned proceeding to provide my expert opinion in support of Singular Computing LLC's Responsive Claim Construction Brief.

3. I am being compensated at an hourly rate of $600 for my time spent working on this matter. My compensation is not in any way dependent on the statements I set forth herein or the outcome of this proceeding.

I. EXPERTISE

A. Research and Professional Experience

4. I have been serving as Professor in Electrical and Computer Engineering (ECE) at Texas A&M University, College Station, Texas, since September 2015. At Texas A&M, I have held the titles of Professor, Associate Professor (between September 2010 and September 2015), and Assistant Professor (between June 2004 and September 2010). My research is conducted in three main areas – computer systems, including computer architecture from the circuits up, algorithm acceleration using GPUs, FPGAs, and custom ICs and VLSI circuits; logic and its applications; and interdisciplinary extensions of the first two areas. A detailed description of each of these research areas is included in my *curriculum vitae* submitted concurrently herewith as **Exhibit A**.

5. From September 2020 to August 2021, I was on sabbatical at Arizona State University and the Air Force Research Laboratory, continuing my past work on several topics including neuromorphic circuits, approximate computing, hardware security and hardware machine learning.

6. From September 2011 to August 2012, I was on sabbatical at the University of Texas at Austin, working with Professor Jacob Abraham on several topics such as genomics, sinusoidal, signal-based data transfer and medical electronics.

7. From January 2000 to May 2004, while working as an Assistant Professor in Electrical and Computer Engineering at the University of Colorado in Boulder, I performed research on VLSI logic design automation, VLSI layout design automation, VLSI design methodologies to address Deep Submicron (DSM) issues such as cross-talk and power along with interdisciplinary extensions. One of the key research areas that I worked on had to do with regularity in Integrated Circuit design using a "layout fabric" which I developed during my Ph.D. research.

8. From August 1993 to December 1999, I worked as a Research Assistant with the CAD group under Professors Robert Brayton and A. Sangiovanni-Vincentelli at the University of California at Berkeley. My research topics included CAD and DSM design, Sets of Pairs of Functions to be Distinguished (SPFDs), Binary Decision Diagrams, Engineering Change, Hierarchical Synthesis and Verification, Model Matching and Combinational Verification, Timing Analysis in the Presence of Cross-talk and Multi-valued Logic Synthesis. My Ph.D. thesis was the first work that indicated that a regular layout "fabric" should be used to alleviate DSM problems like cross-talk, manufacturing issues, delay variation, and signal integrity.

9. From August 1989 to July 1993, I worked as a Design Engineer with Motorola's MC88110 RISC and PowerPC 603 microprocessor groups in Austin, Texas. During that time, I was involved in various design areas from Design for Testability to Digital and Analog Circuit Design and high-level design. I was also independently responsible for the design of the factory test controller for the MC88110 and became familiar with various ad-hoc and structured test

methodologies. I designed digital and analog circuitry, as well as the MC88110's input/output buffers and clock PLL logic. In addition, I performed attendant tasks in a "vertical" VLSI design methodology, including high-level modeling, layout design and verification, full-chip integration, and global and detailed routing.

10. From August 1988 to July 1993, I worked as a Researcher with Professor M. Ray Mercer's group at the University of Texas at Austin, researching topics such as IC testing and Boolean function representation using Canonical XOR-based circuit decompositions.

11. From August 1987 to July 1989, I worked as a Researcher at the University of Texas at Austin. My research included Computer Architecture and Memory Interface design, applied in the context of the METRIC multi-threaded RISC microprocessor which was being developed by Professors Donald Fussell and Roy Jenevein at the time.

B. Education

12. In 1987, I received a Bachelor of Science degree in Electrical Engineering from the Indian Institute of Technology in Kanpur, India. I obtained a GPA of 3.72 and was ranked fourth in my class of sixty students.

13. In 1989, I obtained my Master of Science degree from the Department of Electrical and Computer Engineering at the University of Texas in Austin, Texas. I was awarded the Microelectronics and Computer Development (MCD) Fellowship and maintained a GPA of 3.909.

14. The thesis for my Master's Degree was titled "*The Design of the METRIC Memory Interface and Memory System*" and involved research surrounding the design of the memory interface of METRIC, a multi-threaded RISC Microprocessor.

15. From 1993 to 1999, I attended the University of California at Berkeley, where I obtained my Doctorate Degree from the Department of Electrical Engineering and Computer

Sciences. From 1993 to 1994, I was awarded the California MICRO Fellowship. I maintained a GPA of 3.963 during my Ph.D. studies.

16. My Dissertation was titled "Cross-talk Noise Immune VLSI Design using Regular Layout Fabrics".

C. Publications, Awards, Grants, and Summary

17. I have authored a total of over 263 peer-reviewed publications, including 40 journal papers, 181 conference papers, and 42 workshop papers. Among these papers, five have received a best paper award, while six other papers received best paper nominations. Additional journal papers and conference papers are currently undergoing peer review. In addition, I have co-authored 9 research monographs, 1 edited research monograph, and 3 book chapters.

18. I was invited to serve as a panelist at a conference seven times and have presented two conference tutorials.

19. I am currently a named inventor on six United States patents and one provisional patent application, with two applications currently pending.

20. I have two current research grants totaling $600,000. The total amount of the grants in which I have been involved in to date is $13.66 million, of which $2.59 million is my portion. Some of these grants are with colleagues in the Electrical and Computer Engineering department, as well as other academic departments at Texas A&M University. My research has been funded both through the government and through industrial sources.

21. I currently serve, or have served, as the following:

- Associate Editor, ACM Transactions on Design Automation of Electronic Systems
- Associate Editor, IEEE Transactions on Computers

- Associated Editor, MDPI Journal of Electronics
- EDA Track Co-Chair for ICECS (2014)
- Panel Chair for Texas WISE (2014)
- Track Co-Chair for ICECS (VLSI Systems, Applications and Computer Aided Design Track) (2013)
- Poster Session Chair for Texas WISE (2013)
- Advisory Committee for HotPI (2013)
- Panel Session Chair for SLiP (2013)
- Track Chair (Logic Track) for ICCAD (2009-2010 and 2015-2017)
- Track Chair (Logic Track) for DAC (2016-2017)
- General Chair for IWLS (2009)
- Technical Program Chair for IWLS (2008)
- Track Co-Chair (Computer Aided Network Design (CANDE) Track) for ISCAS (2008-2010)
- Track Co-Chair (Test and Methodologies Track) for ICCD (2007)
- Panel Chair for ITSW (2009)
- Publicity Co-Chair for GLS-VLSI (2009)
- Member of the TPC (several conferences)
- Session Chair (several conferences)

22. I have also received several awards, including the "Outstanding Professor Award" from the Electrical and Computer Engineering Department at Texas A&M University in 2007 and 2020, the "Association of Former Students' Distinguished Achievement Award in Teaching" in 2009, and the "Association of Former Students' College-level Teaching Award" in 2019.

23. I have over 33 years of research and professional experience in the field of Electrical and Computer Engineering, including in the areas of computer systems, logic synthesis and algorithms for VLSI Design and interdisciplinary extensions.

24. I have taught numerous undergraduate and graduate level courses at Texas A&M in the area of Electrical and Computer Engineering. I have graduated twelve Ph.D. students, seventeen M.S. students, and twelve B.S. Honors students. I am currently advising four Ph.D. students, two M.S. students, and three B.S. research students. I have also advised 37 undergraduates, four of which received an award for their research. Eight research papers in international conferences (one invited) have resulted from my work with undergraduates, and the dissertation of one of my Ph.D. students was nominated for the ACM Best Dissertation Award in 2014.

25. On the basis of my education and the experience described above, I am qualified to give the opinions set out herein.

26. In my opinion, a person of ordinary skill in the art would have an undergraduate degree in electrical engineering or equivalent field, which would include a course in statistics.

II. OPINION

27. In my opinion, the term "repeated execution" in the asserted claims of the patents-in-suit (*i.e.*, U.S. Patent Nos. 8,407,273, claim 53; 9,218,156, claim 7; and 10,416,961, claims 4 and 13) is *not* indefinite; indeed, it is readily understandable by a person of ordinary skill in the art ("POSITA") when read in the context of the conjoining claim language, and is explicitly tied to finding "the statistical mean" of the numerical values of the output. For example, claim 53 of the '273 patent recites in relevant part as follows:

> … for at least X=5% of the possible valid inputs to the first operation, ***the statistical mean, over repeated execution of the first operation on each specific input … of the numerical values represented by the first output signal of the LPHDR unit executing the first operation on that input*** differs by at least Y=0.05% from the result of an exact mathematical calculation of the first operation on the numerical values of that same input …

'273 patent, claim 53 (emphasis added); *see also* '156 patent, claim 7; *see also* '961 patent, claims 4, 13.

28. This claim language clearly informs a POSITA that "repeated execution of the first operation" (for example, the multiplication operation 2.0 × 1.0) by the LPHDR unit is used to obtain the "statistical mean" of the numerical values represented by the "first output signal" of the LPHDR unit, as that unit repeatedly executes the first operation. As I explain in more detail below, a POSITA would know that though numerical values generated by the output of any usable computer that is repeatedly executing an operation may slightly fluctuate from execution to execution, the average of those numerical values goes from being an arithmetic average that potentially has an unstable value when computed based on a small number of executions, to a stable value that does not meaningfully fluctuate. This is based on the well known statistical concept called the "Law of Large Numbers", which is taught in a basic statistics course, which a POSITA would have taken in their sophomore or junior year. Based on this elementary knowledge, a POSITA would know that the "statistical mean, over repeated execution", of the "the numerical values represented by the first output signal" would require them to conduct a large enough number of repetitions until the statistical mean reached its stable value.

29. In other words, once the average numerical value of the output for that operation's repeated execution stabilizes, that average *never materially changes again*, even as more and more of the same operation are repeatedly executed. Therefore, in the context of the claimed inventions, a POSITA would understand that "repeated execution" requires the LPHDR unit to

repeatedly execute the "first operation" as many times as necessary to establish that inevitable stable average of the numerical values represented by the unit's output signal, as it repeatedly executes that "first operation". In the field of statistics, an average that represents *all* members of a population is referred to as a population mean. A POSITA would understand that the term "statistical mean", in the context of the asserted claims, refers to the population mean.

30. The purpose served by the claim term "statistical mean" is to characterize the numerical values represented by an output signal (hereinafter, "output values") of a computing device, which typically executes on the order of a billion operations a second.[1] A POSITA would understand that it would be technically erroneous and absurd to generate a "statistical mean" of the output value of a repeated operation, as that term is used in the context of the asserted claims, by repeating an operation just ten times and taking the arithmetic average of the ten output values.

31. Google's argument based on taking an average of 10 or fewer samples is technically erroneous as a matter of engineering as well as common sense. Accordingly, the test results in Google's expert report are unreliable on their face.

32. An in-depth analysis of Google's report does not alter the unreliability of its test results. In his declaration, Dr. Wei worries that a POSITA would be left in the dark as to how many repeated executions should be made before being able to characterize the output value of the claimed device. *See* Wei Decl. at 32. In my opinion, however, this is hardly a problem. A POSITA would only have to apply standard statistical analysis to properly characterize the device's output values.

[1] A POSITA would also understand that limiting the number of repeated executions to ten would serve no practical purpose, as typical computing devices are capable of performing *billions* of floating-point operations every second. *See, e.g.*, Linpack performance Haswell E (Core i7 5960X and 5930K).

33. As the Singular patent specification explains, devices that use analog signals to represent numbers "introduce noise into their computations". *See* '273 patent at 4:12-13. Performing the same operation twice with identical inputs will statistically produce different output values, and initially, a fluctuating arithmetic *average*. However, a POSITA would understand that the output values of repeated executions of the same operation must exhibit the following statistical behavior for the computer to be usable: the average of those output values, over repeated executions, goes from being an arithmetic average that potentially has an unstable value when computed based on a small number of executions, to a stable statistical mean that does not meaningfully fluctuate. Moreover, once enough repeated executions have occurred, that statistical mean no longer materially changes no matter how many more repeated executions are conducted over the useful life of the computer.

34. As shown in the graph below, even using Dr. Wei's experimental conditions as a starting point, a POSITA using fundamental statistical analysis can determine the claimed "statistical mean" in mere microseconds (assuming that the executions are repeated every nanosecond):

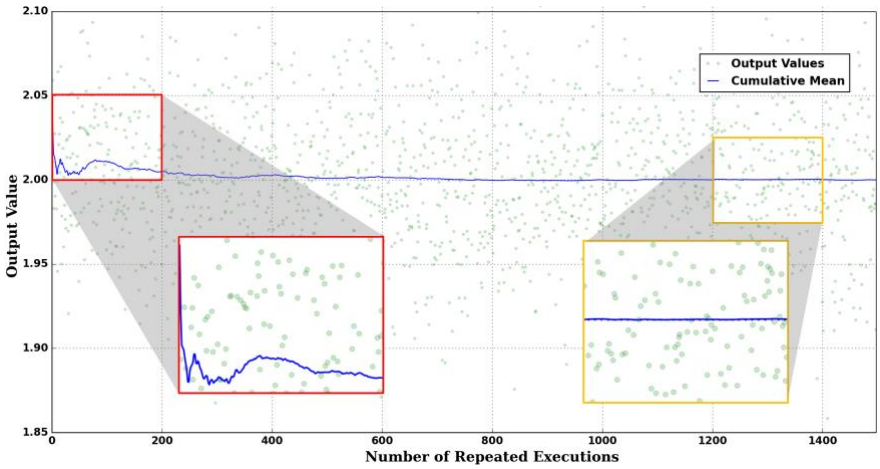

The green dots on the graph above represent a sequence of 1500 "repeated executions" of a single operation (in particular, the multiplication operation 2.0 × 1.0), plotted with the initial

executions on the left and the later executions on the right. The vertical position of each green dot represents the output value of the corresponding execution of the operation[2]. It is apparent that these output values individually fall within a wide range and show no discernible pattern. The blue line shows the average of the output values of the first *x* repeated executions (for example, at *x*=10, the line represents the average of the output values of the first 10 executions). At first, near the left side of the graph, the arithmetic average of the output value is unstable and fluctuates significantly over short periods of time (as shown by the magnified portion of the graph outlined using the **red box**). However, the arithmetic average of the output value begins to stabilize with more repeated executions of that single operation, holding steady at a value of 2.00. Thereafter, no matter how many other repeated executions are added, the arithmetic average of the output value *never varies from 2.00 by more than a few hundredths of a percent* (as shown by the magnified portion of the graph outlined in **orange**). *This portion of the graph represents the "statistical mean, over repeated execution … of the first operation" recited in the claims*. Every execution unit will in short order generate a stable and unchanging average output value. In this duration, a typical execution unit will perform billions of repeated executions. The graph above demonstrates the technical absurdity of Google's argument and the lack of any need to identify the precise number of executions to ascertain the claimed "statistical mean." In short, in a usable computer, repeated execution of a single operation in the context of the asserted patents, will *always* reveal a single statistical mean after a few seconds. This phenomenon is due to the statistical law known as the Law of Large Numbers, as described heretofore.

[2] The values shown in this graph were generated according to exactly the same specifications that Dr. Wei used in generating his ten samples, *i.e.*, each value "differs by up to ±5% from an exact value of 2." Wei Decl. ¶ 38. Although not explicitly stated in Dr. Wei's report, I have assumed that the samples follow a normal distribution with a standard deviation of 0.02.

35. A POSITA would be familiar with the Law of Large Numbers and the other statistical principles described above. Statistics is usually taught within the first couple of years in an undergraduate engineering program, and these concepts are well within the bounds of undergraduate statistics. Indeed, Dr. Wei himself, in his own publications, refers to the concept of "statistical average" in a way that clearly assumes the reader will be familiar with the concept that a mean stabilizes over time. *See*, *e.g.*, Zhang, Brooks, Wei, *et al.*. "A 20μW 10MHz relaxation oscillator with adaptive bias and fast self-calibration in 40nm CMOS for micro-aerial robotics application," p. 435.

36. In its brief, Google refers to a processor that does not exhibit a stable statistical mean due to "heat." *See* Google Br. at 10. A person of ordinary skill in the art would understand that such devices would not satisfy the "repeated execution" limitation of the asserted claims, and that they would not serve any useful purpose as "execution units."

37. Lastly, Google's brief makes the unsupported and unexplained assertion that "the 'first input signal' does not itself have a dynamic range." Google Br. at 19. This assertion is technically incorrect. A signal, such as the first input signal of the claims, can indeed have a dynamic range, that can be computed based on the levels of the signal in question. The well known term in the art, "signal to noise ratio" (or "SNR"), is one example in which the level of a signal is used as part of a ratio, as it would be in a dynamic range calculation.

I declare under penalty of perjury under the laws of the United States of America that the foregoing is true and correct.

Dated: Feb. 8, 2021

Sunil P. Khatri, Ph.D.